

ABSTRACT OF THE DISCLOSURE

An object-based storage may employ a transactional interface and a copy on write
5 protocol. The inode for the file (which is identified by the I-number and indicates one or
more blocks storing the file data) may be copied to a working copy of the inode, and the
working copy may be updated to indicate the newly allocated blocks. In response to a file
commit (e.g. a file close or a file synchronization command), the working copy of the
inode may be atomically written to a non-volatile storage. In this manner, the updates
10 performed during the transaction may be atomically committed to the file. The inode
itself may be stored in a file (referred to as the inode file), and the storage may perform
updates to the inode file atomically as well. The inode file may be modified using the
copy on write protocol described above, and the master inode (mapping the blocks of the
inode file) may be atomically written to non-volatile storage to atomically update the
15 inode file. In one particular embodiment, the storage employs a journal technique for
updates to the inode file.

CONFIDENTIAL